

Introduction to Maximum Likelihood Estimation

S. Gliske, University of Michigan

August 7, 2009

Abstract

Maximum Likelihood Estimation (MLE) is one of the best and most standard methods of density estimation. The purpose of this document is to collect various generalities regarding MLE relevant to HERMES analysis into one document. Specific items regarding interpretation are included in other documents.

1 Differential Yields and Probability Densities

A differential yield over D -dimensions can be defined as members of the set \mathcal{F}_D , the set of all functions that map a subdomain \mathcal{D} of a D -dimensional real space to non-negative finite real numbers. A probability density function (PDF) $p(\mathbf{x})$, $\mathbf{x} \in \mathbb{R}^D$, is a normalized yield,

$$\int_{\mathcal{D}} d^D \mathbf{x} p(\mathbf{x}) = 1, \quad (1)$$

whose range is finite. If p has a parametric form depending on parameters $\boldsymbol{\alpha}$, then we write $p(\mathbf{x}|\boldsymbol{\alpha})$. This is read as the probability density at \mathbf{x} given $\boldsymbol{\alpha}$. Such a quantity is a conditional probability density.

2 Maximum Likelihood Estimation

Given a parametric conditional probability $p(\mathbf{x}|\boldsymbol{\alpha})$ and a data set $\{\mathbf{x}^{(i)}\}_{i=1}^n$, the likelihood of $\boldsymbol{\alpha}$ given the data is

$$L(\boldsymbol{\alpha}) = \prod_{i=1}^n p(\boldsymbol{\alpha}|\mathbf{x}^{(i)}). \quad (2)$$

This can be re-expressed in terms of our chosen $p(\mathbf{x}|\boldsymbol{\alpha})$ by using Bayes rule, yielding

$$L(\boldsymbol{\alpha}) = \prod_{i=1}^n \frac{p(\boldsymbol{\alpha})}{p(\mathbf{x}^{(i)})} p(\boldsymbol{\alpha}|\mathbf{x}^{(i)}). \quad (3)$$

The PDFs $p(\boldsymbol{\alpha})$ and $p(\mathbf{x}^{(i)})$ are called the priors, and they represent the assumptions about $\boldsymbol{\alpha}$ and \mathbf{x} prior considering them jointly, i.e. the distributions independent of each other. As we make no assumptions about $\boldsymbol{\alpha}$ that are not based on the data, and as we assume \mathbf{x} has no further structure than that parametrized by $\boldsymbol{\alpha}$, both of these terms are constant. Thus we have

$$L(\boldsymbol{\alpha}) \propto \prod_{i=1}^n p(\mathbf{x}^{(i)}|\boldsymbol{\alpha}). \quad (4)$$

It is numerically preferable to work with the logarithm of this expression, the log-likelihood

$$LL(\boldsymbol{\alpha}) = \sum_{i=1}^n \ln p(\mathbf{x}^{(i)}|\boldsymbol{\alpha}) + \text{const.} \quad (5)$$

MLE is a method of determining the optimal parameters, given a data set and a parametric form for $p(\mathbf{x}^{(i)}|\boldsymbol{\alpha})$, and is based on the assumption that the optimal parameters are those parameters that are most likely given the data, i.e. that maximize the above likelihood (or log-likelihood) function. This is usually written

$$\hat{\boldsymbol{\alpha}} = \arg_{\boldsymbol{\alpha}} \max \sum_{i=1}^n \ln p(\mathbf{x}^{(i)}|\boldsymbol{\alpha}), \quad (6)$$

where $\hat{\boldsymbol{\alpha}}$ is the MLE estimate of the parameters. Numerically, one can equivalently minimize the negative of the above expression, although in speaking and writing one usually always refers to the log-likelihood and maximizing, rarely the negative log-likelihood and minimizing.

3 MLE with Yields

A differential yield is related to a PDF by dividing by the integral,

$$p(\mathbf{x}) = \left[\int_{\mathcal{D}} d^D \mathbf{x} f(\mathbf{x}) \right]^{-1} f(\mathbf{x}). \quad (7)$$

and similarly with conditional probability densities. The integral $\int_{\mathcal{D}} d^D \mathbf{x} f(\mathbf{x})$ is the expected total yield, i.e. number of events.

This can be substituted into the log-likelihood function. Note the integral does not depend on the data, and so can be taken out of the sum. The MLE estimate with a non-normalized function then has the form

$$\hat{\boldsymbol{\alpha}} = \arg_{\boldsymbol{\alpha}} \max \left[\sum_{i=1}^n \ln f(\mathbf{x}^{(i)}, \boldsymbol{\alpha}) - n \ln \int_{\mathcal{D}} d^D \mathbf{x}' f(\mathbf{x}', \boldsymbol{\alpha}) \right], \quad (8)$$

where the prime has been added to the integration variable \mathbf{x}' to further distinguish it from the data variables $\mathbf{x}^{(i)}$.

The integral can either be computed analytically or numerically. A typical method of numeric integration in several dimensions is called Monte Carlo integration. In this case, one chooses a set $\{\mathbf{x}'^{(j)}\}_{j=0}^m$ which is distributed uniformly at random within the domain \mathcal{D} . The integral is then approximated by

$$\int_{\mathcal{D}} d^D \mathbf{x}' f(\mathbf{x}', \boldsymbol{\alpha}) \approx \frac{V_{\mathcal{D}}}{m} \sum_{j=0}^m f(\mathbf{x}'^{(j)}, \boldsymbol{\alpha}), \quad (9)$$

where $V_{\mathcal{D}}$ is the volume of the domain. As only the logarithm of this integral occurs in the log-likelihood, the factor $V_{\mathcal{D}}/m$ contributes a term constant with respect to $\boldsymbol{\alpha}$ and can be ignored.

Note that if the set $\{\mathbf{x}'^{(j)}\}_{j=0}^m$ is not chosen uniformly at random, but according to some distribution $p'(\mathbf{x})$, then one actually has approximated

$$\frac{1}{m} \sum_{j=0}^m f(\mathbf{x}'^{(j)}, \boldsymbol{\alpha}) \approx \int_{\mathcal{D}} d^D \mathbf{x}' p'(\mathbf{x}) f(\mathbf{x}', \boldsymbol{\alpha}). \quad (10)$$

This fact is used to relate the Kullback-Leibler divergence to MLE, extending the use of MLE to other fields and extending its interpretation. This fact will also be helpful when including the acceptance, Section 5.

4 Uncertainty

A somewhat approachable treatment of uncertainty calculations and confidence levels can be found in Reference [1], mainly based on the book in Reference [2]. The important result in our case is that the covariance matrix C for the parameters α is the negative inverse Hessian matrix for the log-likelihood,

$$(C^{-1})_{i,j} = -\frac{\partial^2}{\partial\alpha_i\partial\alpha_j} LL \Big|_{\alpha=\hat{\alpha}}. \quad (11)$$

The negative can be thought of as arising from the fact that one is maximizing instead of minimizing. Alternately, one can state the covariance matrix is the Hessian matrix of the negative log-likelihood.

5 Acceptance

In any real experiment, both acceptance and smearing affects are present. Thus rather than having simply the true distribution $p_T(\mathbf{x}_T)$, depending on the true values of the variables \mathbf{x}^T , one instead has the measured distribution $p_M(\mathbf{x}_M)$, related by some conditional probability $p(\mathbf{x}_M|\mathbf{x}_T)$, the differential probability of measuring \mathbf{x}_M given that \mathbf{x}_T really occurred. The expression is a Fredholm integral equation,

$$p_M(\mathbf{x}_M) = \int_{\mathcal{D}_T} d\mathbf{x}_T p(\mathbf{x}_M|\mathbf{x}_T) p_T(\mathbf{x}_T), \quad (12)$$

and is further treated in HERMES Internal Note 08-015 [3], among other places [4].

Assuming smearing effects are small is assuming the conditional probability $p(\mathbf{x}_M, \mathbf{x}_T)$ is approximately proportional to $\delta^D(\mathbf{x}_M - \mathbf{x}_T)\epsilon(\mathbf{x}_M)$, or that the measured distribution is

$$p_M(\mathbf{x}) \propto \epsilon(\mathbf{x})p_T(\mathbf{x}). \quad (13)$$

Note the proportional sign, since ϵ changes the normalization.

5.1 Correcting with Normalization Monte Carlo¹

To account for acceptance (neglecting smearing), one can follow the method presented by Andy Miller [6], repeated here with slight typographical changes. Factor $p_M(\mathbf{x})$ into the angular integrated unpolarized cross-section σ_{UU} and remaining portion, g . Let \mathbf{y} denote the non-angular variables and ϕ the angular variables. One would then like to use the non-normalized function f ,

$$f(\mathbf{y}, \phi, \alpha) = \epsilon(\mathbf{y}, \phi)\sigma_{UU}(\mathbf{y})g(\mathbf{y}, \phi, \alpha) \quad (14)$$

to find the MLE estimate of the parameters α . Note, in taking the logarithm of f , the factors $\epsilon(\mathbf{y}, \phi)$ just contribute a constant offset to the log-likelihood and thus can be ignored. However, this is not true for the normalization integral,

$$\int_{\mathcal{D}} d\mathbf{y}d\mathbf{x} f(\mathbf{y}, \phi, \alpha) = \int_{\mathcal{D}} d\mathbf{y}d\mathbf{x} \epsilon(\mathbf{y}, \phi)\sigma_{UU}(\mathbf{y}) g(\mathbf{y}, \phi, \alpha), \quad (15)$$

as the logarithm is taken after the integral.

¹This section describes the rational behind this method. The applicability and interpretation of this method is discussed in Reference [5]. Do not use this method until reading this other document.

Although the product $\epsilon(\mathbf{y}, \phi) \sigma_{UU}(\mathbf{y})$ is unknown, Monte Carlo data can be generated according to this distribution. Thus using Monte Carlo integration and Monte Carlo data, one has

$$\int_{\mathcal{D}} d\mathbf{y} d\mathbf{x} f(\mathbf{y}, \phi, \alpha) \approx \frac{1}{m} \sum_{j=1}^m g(\mathbf{y}^{(j)}, \phi^{(j)}, \alpha), \quad (16)$$

where m is the number of Monte Carlo data points used. Putting this into the MLE expression, one finally has

$$\hat{\alpha} = \arg_{\alpha} \max \left[\sum_{i=1}^n \ln g(\mathbf{y}^{(i)}, \phi^{(i)}, \alpha) - n \ln \sum_{j=1}^m g(\mathbf{y}^{(j)}, \phi^{(j)}, \alpha) \right], \quad (17)$$

recalling the sum over i is using real data, while the sum over j is using Monte Carlo data. If, instead, one mistakenly uses the real data in both sums, the result approximates

$$\sum_{i=1}^m g(\mathbf{y}^{(i)}, \phi^{(i)}, \alpha) \approx m \left[\int_{\mathcal{D}} d\mathbf{y} d\phi g^2(\mathbf{y}, \phi, \alpha) \right] \left[\int_{\mathcal{D}} d\mathbf{y} d\phi g(\mathbf{y}, \phi, \alpha) \right]^{-1}, \quad (18)$$

a quantity not useful in this context, rather than approximating the needed normalization integral.

References

- [1] Geyer, C.J. Class Notes. <http://www.stat.umn.edu/geyer/5102/notes/fish.pdf>
- [2] Degroot, M. and Schervish, M. *Probability and Statistics*, Addison-Wesley Publishing, 2002 (ISBN: 0201524880).
- [3] HERMES Internal Note 08-015 discusses the mathematics of unfolding and presents unfolding methods using histograms and KDEs.
- [4] Additonal references and current information is available at http://hermes-wiki.desy.de/index.php/Unfolding_of_Detector_Smearing_and_QED_Radiative_Effects
- [5] Gliske, S. “Note Regarding MLE and Asymmetries.” http://www-hermes.desy.de/LOTS/0908/sgliske/Gliske.MLE_Interpretation.pdf
- [6] Andy Miller has discussed the matter on several occasions, including http://www-hermes.desy.de/groups/mgmtgrp/COLLABMEETINGS/TRANSVERSITY_JUNE06/MaxLike.pdf
http://www-hermes.desy.de/groups/mgmtgrp/COLLABMEETINGS/OCT05/TRANSVERSITY/Andy_unbinned.pdf
http://www-hermes.desy.de/groups/mgmtgrp/COLLABMEETINGS/DEC05/TRANSVERSITY/Andy_MaxLike.pdf